Subject: **RECOMMENDATIONS FOR CHANGES IN GLITE AUTHORIZATION**

Author(s): **Christoph Witzig**

Distribution **JRA1 management, TCG**

1. INTRODUCTION

In this note we put forward a set of recommendations for changes in the gLite authorization mechanisms. These recommendations were derived during the MJRA1.7 end-to-end study on authorization mechanisms in gLite.

The TCG has set the following priorities for deriving these recommendations:

1. The main focus should be on fixing limitations of the current authorization framework
2. New features should only be introduced to the extent that they are needed by
   a. The experiments / VOs
   b. The sites / SAx
   c. JRA1
3. Interoperability with other grid middlewares
4. Use of standards if possible.

In agreement with these recommendations as well as given the available manpower within EGEE-III, we focused rather on improving and gently extending the current authorization framework than proposing a new, radically different authorization solution.

For every recommendation in this document we give

- Some background information
- The recommendation itself
- A recommended timeframe, in which this recommendation should be implemented (six month, one year, 1 ½ years)
- A few examples, where applicable
- A few comments, where applicable.

The recommendations are grouped into the following categories:

1. Recommendations for FQAN pattern matching rules
2. Recommendations for user interface (UI)
3. Recommendations for compute element (CE)
4. Recommendations for workload manager (WMS)
5. Recommendations for the new authorization service

6. Recommendations for data management

Finally, it should be stressed that these recommendation reflect the impressions of the author, which he obtained during the end-to-end study of authorization mechanisms for the gLite middleware. As such they also reflect to a certain extent personal judgement and preferences.

2. RECOMMENDATIONS FOR FQAN PATTERN MATCHING RULES

RECOMMENDATION 1: USE OF STANDARD LIBRARY FOR FQAN PATTERN MATCHING RULES:

Background: Today, we cannot be sure that different services use the pattern matching rules in a consistent way.

Recommendation: A standard library and a test suite should be developed, which implements the FQAN pattern matching rules. They should become part of the standard gLite distribution. Existing code should be modified to use this library wherever possible. Where this is not possible, the existing implementation should be tested against the test suite.

Timeframe: within 6 months

Comment: A first version of such a library has already been implemented.

RECOMMENDATION 2: USE OF WILDCHARS IN FQAN PATTERN MATCHING:

Background: Current pattern matching rules are rather elaborate and difficult to understand. They should be simplified. In addition, it is not clear whether their complexity correspond to a real need.

Recommendation: The only supported wildchar should be the asterix character (“*”), and it should only be used

- in the group string after the trailing slash (“/”) or
- the “role=” string to denote all possible roles

Examples:

- /VO1/analysis/*: matches all subgroups of /VO1/analysis, but not the parent group (/VO1/analysis)
- /VO1/analysis/*/role=production: matches all subgroups of /VO1/analysis in the production role
- /VO1/analysis/*/role=*: all subgroups of /VO1/analysis where the role parameter has been set.
- /VO1/*/higgs: invalid expression (wildchar in the middle)
- /VO1/analysis*: invalid expression: wildchar not after trailing slash of the group
- /VO1/analysis/?iggs: invalid expression: invalid wildchar at an invalid location

Timeframe: within 6 months

Comments:

1. The use of the role parameter is not changed by this recommendation. I.e. the group and role section of the primary FQAN is still matched independently.
2. It was considered to abandon the use of wildchars at all. However, this was deemed to drastic a change, as wildchars fulfil very useful roles in blacklists and LCMAPS.
3. It was considered to abandon the existing scheme of groups and roles in favour of having only groups or roles. However, it was decided to keep the existing model, as it has been widely accepted and is understood by the majority of the grid administrators. In addition, changing the group and role concept would lead to interoperability issues with other grid middlewares.
4. The use of "wildwords" was considered, but rejected. Its main benefit is that it leads simpler, clearer rule. However, it was felt that the benefit of simpler rules does not outweigh the rather drastic change of established pattern matching algorithms.
5. Note: the rule for an entire VO still consists of two patterns
   - /VO1/* : to cover all subgroups
   - /VO1 : to cover the VO itself

RECOMMENDATION 3: DEVELOPMENT OF A FQAN DEBUGGING TOOL:

Background: Currently the primary FQAN is used by several packages in the WMS and CE.

Recommendation: A command line tool, available on the WMS and CE, should print the authorization decisions for every package in the CE and WMS. The input parameter of this debugging utility can either be the primary FQAN or a proxy certificate.

Timeframe: within six months.

Examples:

1. On the WMS the command line tool would print out all VOViews and their ACBRs, that are considered for this primary FQAN.
2. On the CE the command line tool would print the decisions of LCAS, LCMAPS as well as the share to which this primary FQAN is assigned.

3. RECOMMENDATION FOR USER INTERFACE

RECOMMENDATION 4: THE USER SHALL BE ABLE TO SPECIFY THE FQAN USED IN THE JOB SUBMISSION:

Background: Currently the primary FQAN has the special role in the job submission as well as the data management. In the job submission it is being used to map the user at the CE to a given primary Unix GID. In the data management it is used to denote group ownership of files created.

Recommendation: We recommend that the user shall be able to specify the FQAN to be used in the job submission either as a parameter in the JDL file or as a variable in scripts.

Timeframe: within one year

Comments:

1. Adding a directive for choosing the FQAN in the JDL specification allows supporting this mechanism in the standard gLite job submission, either through the WMS or through direct invocation of the CE. The use of a variable in scripts allows the use of this functionality in pilot jobs.
2. The affected services are the WMS and CE (LCAS/LCMAPS).
3. This modification supports the use-case that a job gets executed under a Unix GID, which is independent of the FQAN used for data management.
4. This recommendation is a very limited option of giving the user the choice to specify, which FQAN to consider during an operation on the grid. In its full form one could envisage that the user specifies one FQAN as primary FQAN for the job submission, another FQAN as primary FQAN for creating output files and yet another primary FQANs for different manipulations of file catalogues. These use cases are all valid. However, it was felt that the benefit of implementing a general solution does not justify the substantial effort that such a change would require - particularly given the limited means of JRA1 in EGEE-III. E.g. for the data management such a modification would not only mean changing in the entire stack from the top level interface down to file access, but also a change in the standardized SRM interface, i.e. the standardization body would have to be involved as well.
5. One could envisage a scheme in which additional information is embedded in the proxy to indicate to services, which FQAN to take for a given operation. However, we recommend not "abusing" the proxy as carrier for all kind of job specific information. The proxy certificate is a credential and should remain so.
6. In the long run we recommend implementing a mechanism, which allows using different FQANs for different operations on the grid. Such an effort should be coordinated with other grid projects and standardization bodies.

4. RECOMMENDATIONS FOR THE CE

RECOMMENDATION 5: LCMAPS SHOULD TAKE THE MOST SPECIFIC MATCH INSTEAD OF THE FIRST MATCH

Background: While finding the Unix GID for the primary FQAN, LCMAPS currently uses the first match. Thus the order of the patterns in the LCMAPS configuration files matters in the determination of the GID. It is also one possible cause for inconsistencies between authorization in the CE and the match making of the WMS.

Recommendation: Instead of taking the first match, LCMAPS should return the GID corresponding to the most specific match. The most specific match is the match, which matches the most characters excluding any possible wildchar.

Timeframe: within six months

Examples: Assume the primary FQAN is /VO1/analysis/higgs.

1. For the patterns /VO1/analysis, /VO1/analysis/higgs, /VO1/analysis/*, /VO1/analysis/higgs/* the most specific match is /VO1/analysis/higgs
2. For the patterns /VO1/analysis, /VO1/analysis/*: most specific match is /VO1/analysis/*

Comments:

1. This change guarantees that the match as determined by LCMAPS is independent of the order in which the patterns as listed in the LCMAPS configuration file.
2. The use of the role parameter is not changed by this recommendation. I.e. the group and role substrings of the primary FQAN are still matched independently.
3. This change allows the WMS matchmaking to determine the LCMAPS mapping through the use of the VOView ACBRs without any further information. (See also Recommendation 9: The WMS matchmaking should only consider one VOVIEW PER CE).

RECOMMENDATION 6: USE OF GJAF IN THE CE

Background: The CREAM CE currently uses two authorization frameworks: gJAF for authentication and authorization decisions in java code and LCAS/LCMAPS within glexec.

Recommendation: The gJAF framework should be abandoned and replaced with a simple authentication check of the certificate and a simple call-out mechanism to the new site authorization service.

Reason: The use of two authorization frameworks in the same service (i.e. the CE) is not justified and may lead to inconsistent authorization decisions.

Comments:

1. gJAF will no longer be supported in EGEE-III.
2. If it turns out that a richer functionality than a minimal authorization call-out is needed at the CE, then the Globus authorization framework should be considered as a solution. It is independent of the Globus code and its continued maintenance seems to be better guaranteed than gJAF.

RECOMMENDATION 7: VOVIEWS SHOULD PUBLISH ALL ACBR INFORMATION RELEVANT FOR MATCH-MAKING IN THE WMS

Background: The VOViews contain a field ACBR, which lists all FQANs that the underlying share supports for a given VO. In the past the information in this field did not properly reflect the configuration of the batch system relevant for the authorization decisions, i.e. the mapping between VOView ABCRs and FQANs..

Recommendation: The GIP and IS must make sure that all ACBR information relevant for the match-making in the WMS is published in the IS. This means that all the information for assigning FQANs to shares has to be published in the IS, including possible wildchars. Otherwise the WMS cannot do the matchmaking properly.

Timeframe: within six months.

Comment: This is a requirement for several components:

1. The GIP, which publishes the VOViews
2. The IS and Glue schema, which must support wildchars

RECOMMENDATION 8: DECOUPLING OF FQAN'S AND SHARES AT THE CE

Background: The CEs currently deployed in EGEE determine the Unix GID based on the primary FQAN. In addition, the Unix GID is used to assign a job to a given share of the batch system. Thus the authorization data is tied to the shares, and the site has to be reconfigured whenever the assignment between FQAN and shares has to be changed, i.e. whenever the VO wants to assign different CPU shares to different groups of jobs.

Recommendation: The tight coupling between FQANs and shares at the CE should be replaced with a mechanism that allows the VO assign different types of jobs to different shares without site reconfiguration.

Timeframe: one year

Comments: A possible solution is proposed in the document
“recommendation_fqan_shares_v1.0.doc” (see https://edms.cern.ch/document/887174/1)

5. RECOMMENDATIONS FOR THE WMS

RECOMMENDATION 9: THE WMS MATCHMAKING SHOULD ONLY CONSIDER ONE VOVIEW PER CE

Background: Currently the matchmaking considers all VOViews, whose ACBRs match the primary FQAN of the proxy certificate, which was submitted together with the job. LCMAPS on the other hand simply takes the first match. This can lead to inconsistencies between the matchmaking and the handling of the job in the CE.

Recommendation: The WMS should only consider one VOView per CE for a given primary FQAN. The selected VOView should be the one, whose ACBR is the most specific match for the primary FQAN.

Timeframe: within six months

Comments:

1. The most specific match is defined in “Recommendation 5: LCMAPS should take the most specific match instead of the first match”. Care has to be taken that this definition for the LCMAPS and the CE always return the same match.
2. This recommendation together with recommendation 5 removes any inconsistencies between the matchmaking and LCMAPS.
3. This recommendation together with recommendation 5 makes sure that neither FQAN ordering information nor DENY tags have to be added to the VOView. Thus, the current glue schema of the VOView does not need to be changed.
4. It has to be verified that this change does not introduce to big an overhead in the matchmaking process. If the overhead on the matchmaking turns out too high, then comment 6 may be the preferred solution.
5. It should be noted that this recommendation might lead to the fact that under certain circumstances not the best VOView is taken for a given job. The following example may explain this situation: Let’s assume that a CE has two shares, whose ACBRs match the

primary FQAN. Let's also assume that the share with the most specific match has the higher load than the share with the less specific match. In this case the "better" share in *that* CE is not considered. However, we feel that the advantages of this recommendation outweigh this particular disadvantage, namely:

a. There are typically many CEs on the grid, where a given FQAN is accepted. Therefore it is rather unlikely that the best VOView is the less specific matched VOView on that very CE *unless* the number of available CEs is small due to requirements on the data availability.

b. The administrator of the CE decided to assign that particular FQAN to a given share. This recommendation guarantees that the site configuration determines the outcome of the mapping decision (local site configuration rules should always come *first*).

6. It is worth considering the alternative to this recommendation: In this case the WMS selects the share and must pass this information to the CE. LCMAPS would then get as additional input parameter the assigned share. If the handling of the ACBRs in the VOViews is correct, then LCMAPS is able to accept this assignment. Otherwise LCMAPS would have to return an error in order to avoid taking an inconsistent decision. While such a mechanism is feasible, we believe that the proposed solution has the advantage that fewer components have to be changed.

6. RECOMMENDATIONS FOR THE NEW AUTHORIZATION SERVICE

RECOMMENDATION 10:

Background: The GP-Box is a service that allows a consistent handling of the authorization information between the CE and the WMS. However, it is currently not deployed within EGEE as part of the gLite distribution. Thus, the question rises, whether the GP-Box should be fully integrated into gLite.

Recommendation: We do not recommend the use of the GP-Box in its current form in gLite for the following reasons:

1. The GP-Box allows the VO to define the user mapping at a remote site and download the corresponding (XACML) policy file to the local site. The local site administrator has then the choice to either accept or deny the new policy. We believe that copying policy files between different administrative domains tends to rather create problems than solve them. A strict division of responsibilities between VOs and sites should be maintained, particularly if the user mapping is involved.

2. The GP-Box provides a graphical user, which allows to edit the policies and co-ordinate them between different components of the grid (WMS and CE). This approach is very suited for a small grid or campus grid infrastructure, where the same person administrates all components and very few components are involved. However, we believe this approach is not suited in a large, heterogeneous environment such as EGEE.

RECOMMENDATION 11:

Background: JRA1 has been given the task to design and implement a new authorization service that supports XACML policies. The following EGEE-III partners have been assigned to this task: CNAF, HIP, NIKHEF and SWITCH. The design of this service is a part of this authorization work and will be described in another document.

Recommendation: One component of the new authorization service is a policy decision point (PDP) that understands XACML policies. The design of this authorization service should take the existing GP-BOX code base into account and re-use or re-engineer components on an as needed basis.

Comment: This detailed design of this new authorization service is currently in progress.

7. RECOMMENDATIONS FOR THE DATA MANAGEMENT

RECOMMENDATION 12: ADOPTION OF THE DPM SECURITY MODEL BY OTHER STORAGE SOLUTIONS

Background: The DPM storage element (SE) has implemented an authorization mechanism based on the use of the FQANs in the user’s proxy certificate. Other SE implementations have not adopted this mechanism.

Recommendation: All storage element implementations should adopt the DPM authorization model.

Comment: This decision has already been taken. This recommendation only reinforces the chosen path.